Draft version February 22, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1602.06286v1 [astro-ph.SR] 19 Feb 2016

CONSTRAINING WHITE DWARF STRUCTURE AND NEUTRINO PHYSICS IN 47 TUCANAE

R. Goldsbury[1], J. Heyl[1], H. B. Richer[1], J. S. Kalirai[2,3], P.E. Tremblay[4]

Draft version February 22, 2016

ABSTRACT

We present a robust statistical analysis of the white dwarf cooling sequence in 47 Tucanae. We combine HST UV and optical data in the core of the cluster, Modules for Experiments in Stellar Evolution (MESA) white dwarf cooling models, white dwarf atmosphere models, artificial star tests, and a Markov Chain Monte Carlo (MCMC) sampling method to fit white dwarf cooling models to our data directly. We use a technique known as the unbinned maximum likelihood to fit these models to our data without binning. We use these data to constrain neutrino production and the thickness of the hydrogen layer in these white dwarfs. The data prefer thicker hydrogen layers ($q_{\rm H} = 3.2 \times 10^{-5}$) and we can strongly rule out thin layers ($q_{\rm H} = 10^{-6}$). The neutrino rates currently in the models are consistent with the data. This analysis does not provide a constraint on the number of neutrino species.

1. INTRODUCTION

White dwarfs are often utilized as a way to indirectly test physical properties that are not directly observable themselves. White dwarfs have been used to measure the distances (Renzini et al. 1996; Zoccali et al. 2001; Woodley et al. 2012), and ages (Hansen et al. 2007, 2013) of globular clusters. They have been used to constrain the initial mass function and age of the galaxy (Wood 1992). The physics of crystallization of dense white dwarf interiors has been indirectly constrained by Winget et al. (2009). The white dwarf cooling sequence has even been used to measure dynamical relaxation in 47 Tucanae (Heyl et al. 2015). White dwarfs in globular clusters have proven particularly useful for this, because all of the objects in a cluster come from an initial population with roughly the same age, extinction, distance, and typically the same initial chemical composition.

In this paper we use the most well-populated white dwarf sequence ever measured in a globular cluster (47 Tuc) to indirectly constrain the thickness of the Hydrogen layer, and the rate of neutrino production in the cores of white dwarfs. This will build on our previous work from Goldsbury et al. (2012), in which we constrained white dwarf cooling models from various groups against multi-band HST photometry in 47 Tuc. These models are again compared to the new data set in Section 4.4.

Our dataset is obtained from nine HST orbits that cover the core of 47 Tuc in ultraviolet and visible filters. In total more than 90% of the inner 5 arc-minutes of the cluster are observed, and over 3500 white dwarfs are detected. This dataset is described in detail in Section 2. The analysis discussed in this paper complements the work in Hansen et al. (2015). In that paper, similar white dwarf cooling models were fit to the ACS data shown as the surrounding annulus in Figure 1. These models were fit by binning the data in two dimensions (color and magnitude) and using the χ^2 statistic to address the quality of the fit.

[1] Department of Physics & Astronomy, University of British Columbia, Vancouver, BC, Canada V6T 1Z1
[2] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD, 21218
[3] Center for Astrophysical Sciences, Johns Hopkins University, Baltimore, MD, 21218
[4] Department of Physics, University of Warwick, Coventry, CV4 7AL, UK

Table 1
Exposure times in seconds, for each of the nine visits

Camera	Filter	Exposure 1	Exposure 2
WFC3	F225W	380	700
	F336W	720	485
ACS	F435W	290	690
	F555W	660	360

Our analysis combines data in separate regions from both WFC3 and ACS images. To take full advantage of the constraining power of this dataset, we use an unbinned maximum likelihood method. The data occupy a three dimensional parameter space comprised of magnitudes in two filters and a radial coordinate in the observed field. Theoretical models are transformed into this space using the measured properties of the detection instrument and the reduction procedures. The data are modified as little as possible and our assumptions about detection efficiency and photometric error are built directly into the model itself. This process is described in detail in Section 3.

2. DATA

Our dataset results from ten total HST orbits during Cycle 21 (GO-12971, PI: Richer), one of which was rejected due to loss of guide stars. Observations were taken between November 2012 and September 2013. In each orbit, the cluster was simultaneously observed using the Wide Field Camera 3 (WFC3) and the Advanced Camera for Surveys (ACS). All of the observed fields are shown in Figure 1. Each orbit was split into two exposures for each of the four filters. These are listed in Table 1. Splitting each observation in two allows us to eliminate cosmic ray strikes. The exposure lengths are chosen to fit nicely around the required buffer memory dumps from the telescope. The observations are designed to take advantage of the maximum possible observation time during each orbit. The split observation also allows for a dither between each exposure.

The raw WFC3 images were combined by drizzling

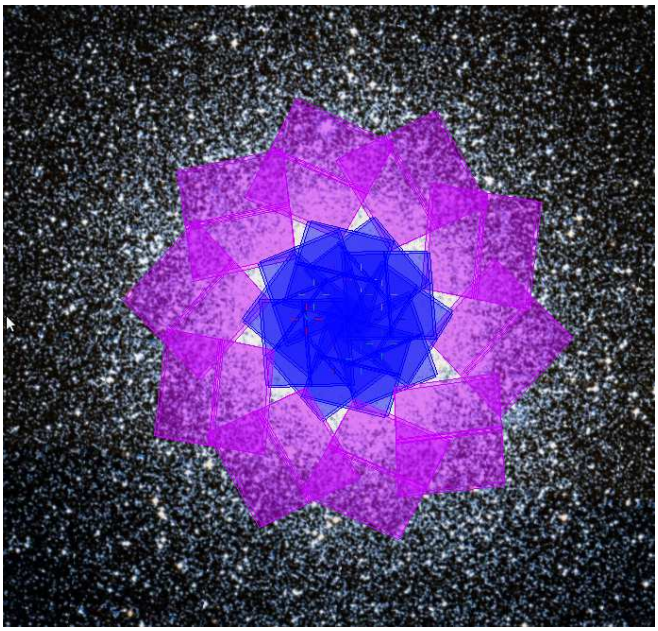

Figure 1. Above are the pattern of the observed fields in the core of 47 Tuc. Fields imaged by WFC3 are shown in blue. Those observed by ACS are in pink. In the final dataset only nine of the ten orientations for both WFC3 and ACS were used due to loss of guide stars during one exposure.

onto a single reference frame (Fruchter et al. 2009). The ACS images were reduced to photometry independently. In total 107,000 objects were detected in the WFC3 fields, and 228,000 objects were detected in the ACS fields. The CMDs resulting from these catalogs are shown in Figure 2.

3. ANALYSIS

3.1. *Artificial Star Tests*

The process by which we generate and measure artificial stars is one of many steps that are integral to the translation of the white dwarf cooling model from theory-space into data-space. By "theory-space" we mean the quantities that come out of the theoretical white dwarf calculations, such as their physical size, their atmospheric composition, their surface temperatures, etc. By "data-space" we mean the dimensions of the space in which the data points exist. The WFC3 data that we are considering are the F225W magnitude, the F336W magnitude, and the radial distance in pixels from the cluster centre of every white dwarf in the sample. So the points in the data set all lie in this three-dimensional space. Similarly, the ACS data-space is made up of F435W, F555W, and radial distance. Any comparison of model to data necessarily requires the model and the data to be transformed to common quantities that are then compared. In our previous paper on white dwarf cooling (Goldsbury et al. 2012) we took an approach that involved modifying the data to produce a time-temperature relation that was then compared to the model. In this paper we do the opposite, which is to leave the data as untouched as possible and perform all of the transformations on the model itself.

The artificial star tests for the WFC3 drizzled image are described below. The process is virtually identical for the ACS data, but is only performed on one of the nine visits. We make the assumption that the cluster's projection is azimuthally symmetric, so the error and completeness measured across one visit are the same for the eight other ACS visits.

First, 1000 values were chosen, evenly spanning the range of 18 to 28 magnitudes in F225W. 1000 values were also chosen, evenly spanning 16 to 26 in F336W. A radial coordinate grid was then defined in the drizzled image frame. This grid consists of concentric rings centred on the core of the cluster with a spacing of 50 pixels between each. Points on each annulus are evenly distributed with a 50 pixel spacing between adjacent artificial stars on the same annulus. This means there are approximately 25,000 points in the drizzled frame. This spacing ensures that the density of artificial stars has no effect on the measurements of the artificial stars. Although the artificial stars can not interfere with measurements of other artificial stars, the real stars can interfere with the measurements of artificial stars, which is what is measured by these tests. The locations of the artificial stars in magnitude and position space are shown in Figure 3.

For each of the 1000 defined magnitudes in each of the two filters, one modified image was generated. Each of these images is identical to the original science image, but with the addition of 25,000 artificial objects, all of the same magnitude. So, in total, 2000 new WFC3 images were generated (1000 for each filter). This process was repeated to generate an additional 2000 new images for the ACS data.

Each of these 2000 images was run through the same reduction pipeline used to generate the catalogue of real objects. Whether or not the artificial object was found during the reduction was recorded. If the object was found, the output photometry was recorded as well.

This output artificial star catalogue is the basis for the error and completeness model. We can index the input artificial F225W magnitudes with i, the input F336W magnitudes with j and the input positions with k. The real catalogue only contains stars that are identified in both filters. Because the positions indexed by k are consistent across both filters, we can consider any value i and j and access the output measurements of those two values at a consistent position k. While measuring the completeness of the artificial stars in this way we assess $i * j * k = 2.5 \times 10^{10}$ unique stars in the data-space.

Considering only stars that are recovered, each unique combination (i, j, k) also has a measured ΔF225W and ΔF336W (the difference between the magnitudes measured by the reduction pipeline and those input), as well as a ΔR. From here on, the difference in input and output position is ignored ($\Delta R = 0$). Our analysis assumes that stars have no uncertainty in their position in the field. All of these data are then used to build a five dimensional array indexed by: $\mathrm{F225W_{input}}$, $\mathrm{F336W_{input}}$, R, ΔF225W, and ΔF336W.

We call this array E, more explicitly: $E(\mathrm{F225W_{input}}, \mathrm{F336W_{input}}, R, \Delta\mathrm{F225W}, \Delta\mathrm{F336W})$. E is generated by counting the artificial stars found after passing through our reduction procedure and dividing the count in each bin by the number of input artificial stars. The completeness of objects at a given pair of magnitudes and position in the field can be

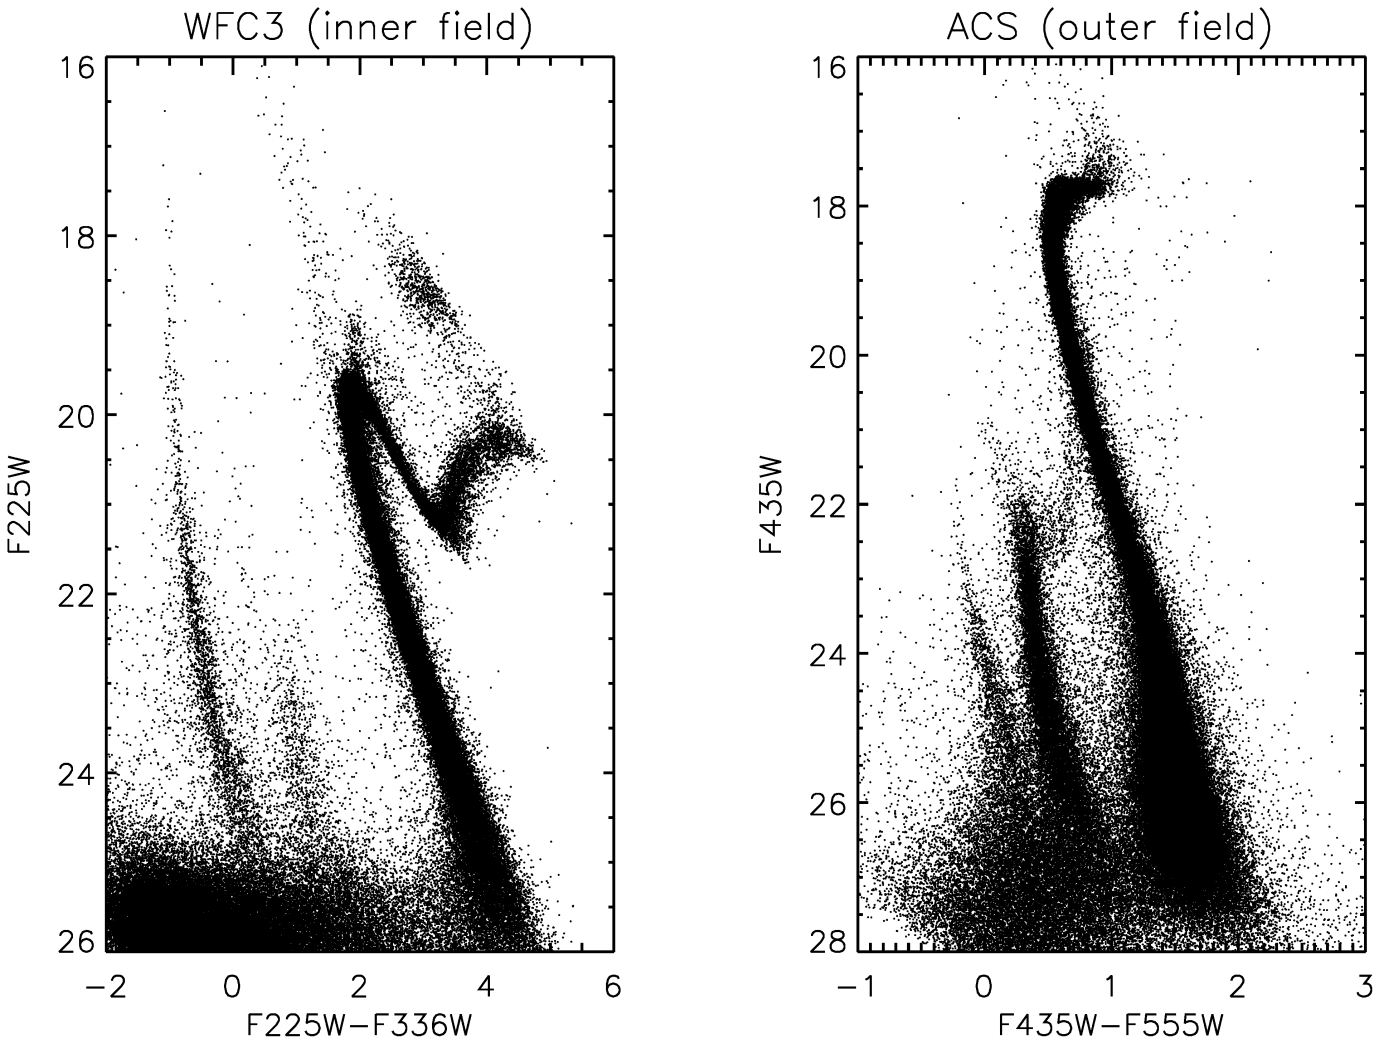

Figure 2. The color magnitude diagrams of the two fields. The central UV data are shown on the left. The visible CMD from the surrounding annulus is shown on the right.

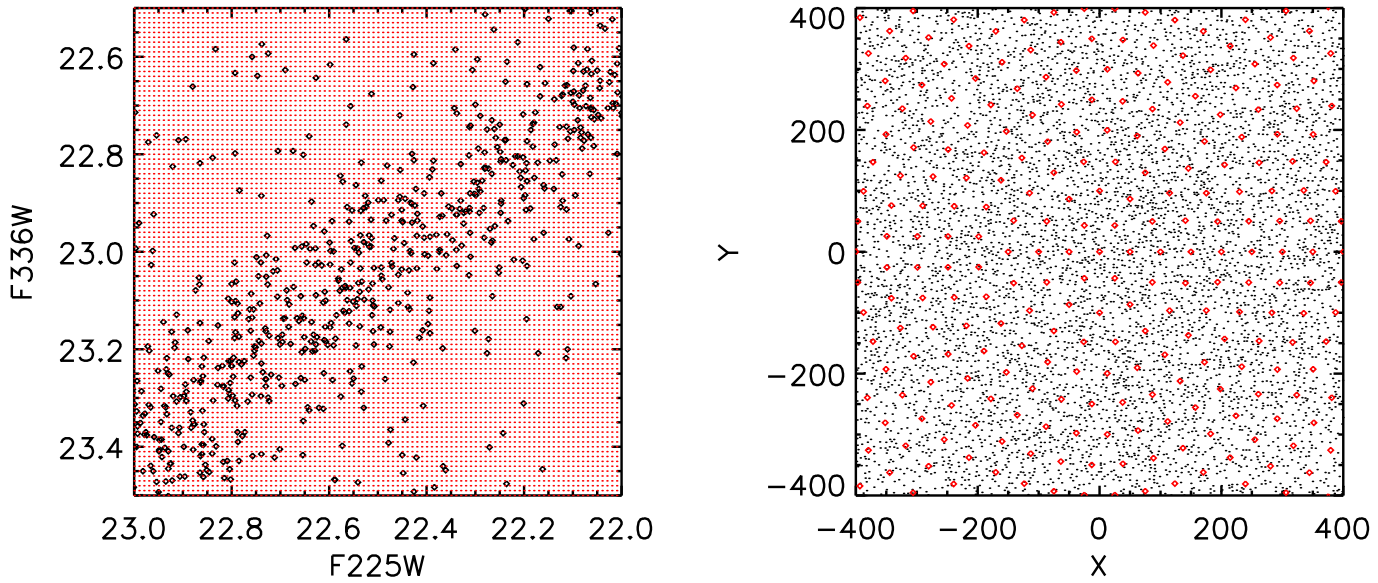

Figure 3. The real data are shown in black and the artificial star input grid is shown in red. The magnitude plot is centred on the white dwarf sequence.

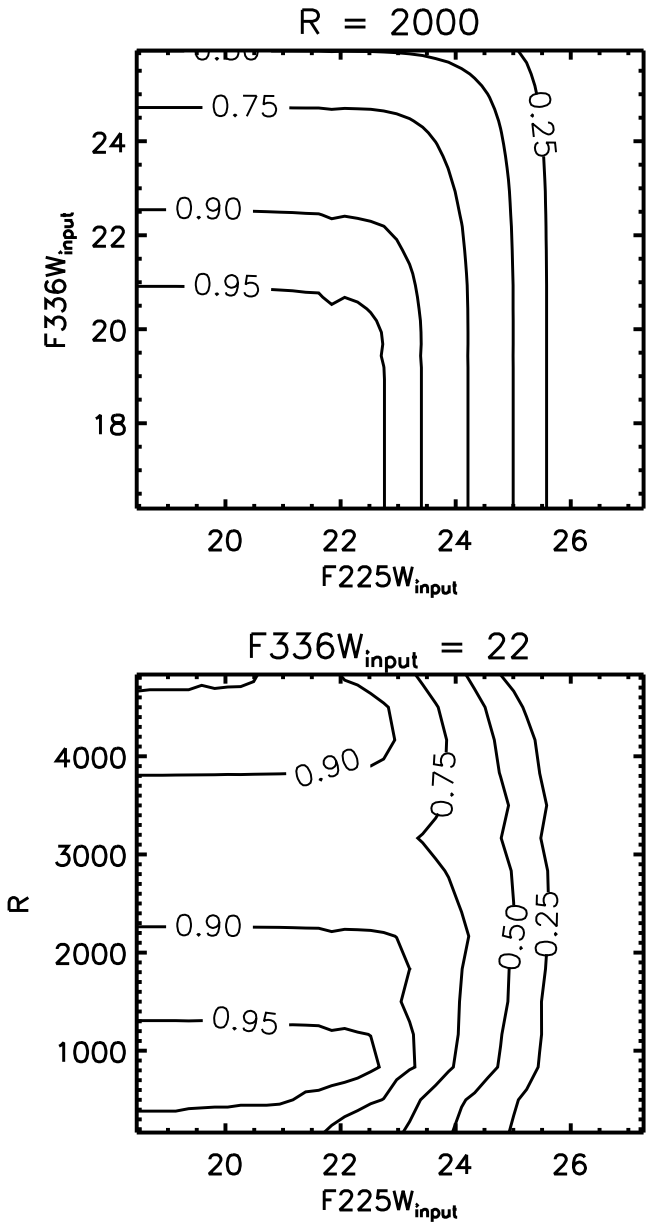

Figure 4. Completeness fraction in magnitude space for a fixed radial distance is shown on the top. Completeness fraction along the two-dimensional plane of R and $\mathrm{F225W_{input}}$ with a fixed value of $\mathrm{F336W_{input}}$ is shown on the bottom. Radial distance is in units of WFC3 pixels, which correspond to 0.04 arc seconds each. These surfaces are the values from C described in Equation 1. C is indexed by three parameters. In each plot above, one of these three is fixed, and C is shown varying across the remaining two.

calculated as

$$C = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E \; d(\Delta\mathrm{F225W}) \; d(\Delta\mathrm{F336W}). \qquad (1)$$

How the completeness varies across the three-dimensional data-space is $C(\mathrm{F225W_{input}}, \mathrm{F336W_{input}}, R)$. The completeness values on two orthogonal slices through this three dimensional space are shown in Figure 4. The odd structure in the bottom plot is real and not an effect of poor sampling. Our observational design leads to much larger total integration time in the core of the cluster than in the surrounding annuli, and this effect contributes positively to completeness as you move toward the core of the cluster. However, stellar density also increases dramatically toward the core, which decreases the completeness. These two effects together lead to the behaviour seen in the bottom figure.

Another way to slice E is to fix the first three parameters. The variations in E over the remaining two parameters (ΔF225W and ΔF336W) are the error distribution of those input stars. Two such distributions are shown in Figure 5.

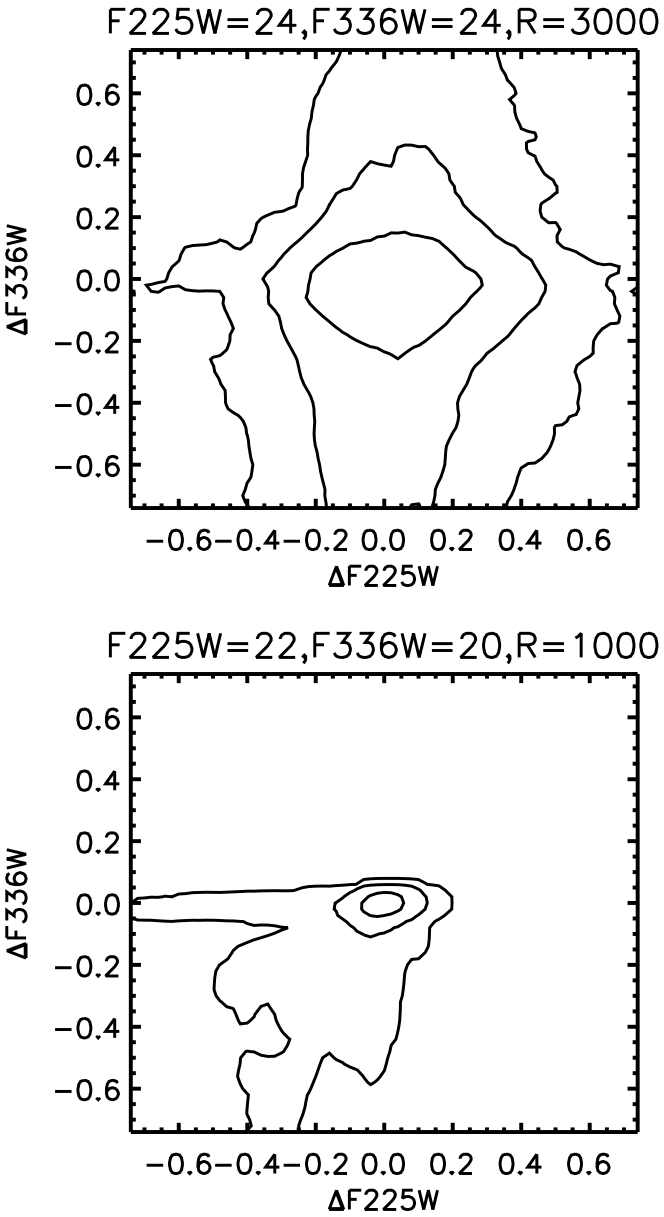

Figure 5. 1, 2, and 3σ contours (enclosing 68%, 95%, and 99.7%) of the photometric error distribution are shown at two points in the data-space. The top panel shows the error distribution for faint stars close to the outside of the inner WFC3 field. The lower panel shows the error distribution for brighter stars closer to the centre. R is again in units of pixels.

The error distribution describes the expected measured differences between the input and output artificial star

photometry. From the figures one can see that the error distribution is highly non-Gaussian, and changes significantly across the space of the data. The five dimensional array E is the key to transforming from predicted model magnitudes to the likelihood of actually measuring objects at various magnitudes in the catalogue. E is an empirically determined model of our photometric scatter and measurement biases caused by the crowding of the field, the instrument and filters used, the observation pattern, and the reduction process. How this information is used in conjunction with our theoretical cooling models is described in Section 3.3.

3.2. *MESA Models*

To calculate the cooling models we will compare to the data, we used Modules for Experiments in Stellar Astrophysics or MESA (Paxton et al. 2011). We used these modules to perform simulations of stellar evolution starting with pre-main-sequence models with a metallicity of Z= 3.3×10^{-3} appropriate for the cluster 47 Tuc. Wind parameters in the RGB and AGB phases of evolution are adjusted to produce white dwarfs with varying hydrogen layer thickness. The parameter adjusted is η. The mass loss rate on the RGB is calculated from Reimers (1975):

$$\dot{M}_{\rm RGB} = 4 \times 10^{-13} \eta \frac{L}{L_\odot} \frac{R}{R_\odot} \frac{M_\odot}{M}. \quad (2)$$

The mass loss rate on the AGB is calculated from Blöcker (1995):

$$\dot{M}_{\rm AGB} = 4.83 \times 10^{-9} \left(\frac{M}{M_\odot}\right)^{-2.1} \left(\frac{L}{L_\odot}\right)^{2.7} \dot{M}_{\rm RGB}. \quad (3)$$

Hydrogen layers with $q_{\rm H}$ between 10^{-6} and 10^{-3} are produced, where $q_{\rm H}$ is the fraction of the WD mass that is in the hydrogen layer. The thickness of the residual helium layer is $-1.45 < \log_{10}(q_{He}) < -1.30$ for all of the models that reach the white dwarf stage. Since we adjust the wind parameters and not the thickness of these layers directly, the masses of hydrogen and helium vary coincidentally as η is changed. This process is stochastic, and very similar wind parameters can produce white dwarfs with hydrogen layers that differ by orders of magnitude. However, using this revision of MESA, it was not possible to produce helium layers that vary over such a range. The mass of the helium layer stays relatively constant over all of our simulations. As a result, we do not consider the helium layer mass as a separate parameter in our fitting. The fit distributions that we find can be considered as the result after marginalizing over this parameter.

During the evolution of these models, neutrino production is multiplied by a factor A_ν ranging from 0.1 to 3.18. We use SVN revision 5456 of MESA and start with the model `1M_pre_ms_t_wd` in the test suite. We changed the parameters `initial_mass` and `initial_z` of the star and adjusted the parameter `log_L_lower_limit` to -6 so the simulation would run well into the WD cooling regime. We defined the starting time of our cooling models to be the last local maximum in luminosity. After the time we define as t_0, the models only decrease in brightness.

Initially it was thought that the mass of the white dwarf would be a dominant parameter. In fact, using the MESA framework, it is rather difficult to produce white dwarfs that differ significantly from the canonical globular cluster value of $0.53 M_\odot$ by adjusting η. In order to change this mass by a large amount, the main sequence evolution time scale or the turn-off mass need to be adjusted to values that are entirely inconsistent with other observations in 47 Tuc and other clusters. Additionally, even when these models were produced, it was found that varying the mass does not produce a large effect on the predicted distribution of white dwarfs on our CMD. This is especially true when uncertainties in distance are considered, which are largely degenerate with uncertainties in white dwarf mass as seen on the CMD. For this reason, we have chosen to fix this parameter and instead look to the hydrogen layer thickness and the neutrino production as our fit parameters. Both of these have a much larger effect on the white dwarf cooling distribution than the mass does.

When producing these models, we found that most values of η lead to white dwarfs with thicker atmospheres ($q_H > 10^{-4}$), and that the mass of the resulting hydrogen layer is independent of η for most values of that parameter. Althaus et al. (2015) find similarly that the mass of H with which the white dwarf enters its cooling track is not altered by the occurrence of mass loss during these stages. The values of η which are able to produce very thin atmospheres constitute a small portion of those considered. So, we produce a large group of models with varying η and A_ν values. The white dwarfs that result from these models are then reparametrized over a grid of q_H and A_ν. The model grid in this space that we end up fitting to our data does not uniformly represent η. The intention is to provide good coverage over q_H and A_ν, since these are the parameters we are interested in studying. The cooling models in this parameter space are the starting point for the approach described in the next section.

3.3. *Moving the Model to Data-Space*

The result of these calculations is a grid of white dwarf cooling sequences over varying $q_{\rm H}$ and A_ν. Each combination of hydrogen layer thickness and neutrino factor produces one cooling sequence that describes how the physical properties of the white dwarf change with time. The ones we are interested in are effective surface temperature T, surface gravity g, and radius r. These three parameters can be thought of as functions of time t but also $q_{\rm H}$ and A_ν. The state of the model resulting directly from calculations in MESA is

$$T(t, q_{\rm H}, A_\nu), \quad (4a)$$

$$g(t, q_{\rm H}, A_\nu), \quad (4b)$$

$$r(t, q_{\rm H}, A_\nu). \quad (4c)$$

These parameters exist in theory-space and need to be transformed to quantities in data-space. The first step is to determine how the model objects look through the HST filters. To do this we use spectral models for DA white dwarfs described in Tremblay et al. (2011) and the publicly available HST filter throughputs available from

STScI (2009). We also model the reddening using the interstellar reddening model from Fitzpatrick (1999).

The spectral model grid provides the spectral flux density (energy per unit time per unit wavelength), F, at the surface of the object as a function of T and g. Explicitly we can write this as $F(\lambda, T, g)$.

The Fitzpatrick reddening curve is parametrized by $E(B-V)$, and R_V. These are defined as follows. A_V is the extinction in V and A_B is the extinction in B (the fraction of power lost in filter V and filter B respectively). $E(B-V) = A_B - A_V$; and $R_V = \frac{A_V}{A_B - A_V}$. The symbols B and V refer to magnitudes in the B and V filters (centred at 445nm and 551nm). We fix $R_V = 3.1$ which is found to apply in most galactic environments (Fitzpatrick 1999). This value also gives good agreement between our white dwarf model sequence and our data, suggesting that it is appropriate for 47 Tuc as well. We do not vary R_V, and it remains fixed at the standard value of 3.1 for the remainder of our analysis. The reddening curve is called $k(\lambda, e)$, and is the fraction of the power transmitted as a function of wavelength between the object and the observer.

The filter throughput curves describe the fractional power transmittance as a function of wavelength for a given filter. This will be labeled as $S_i(\lambda)$ where i is an arbitrary index to indicate which filter is being used.

All of these objects are then combined as shown in Equation 5 to generate a predicted magnitude:

$$m'_i(T, g, r, e, d) = -2.5\log_{10}\left(\frac{\int_0^\infty \lambda F S_i k d\lambda}{\int_0^\infty \lambda F_0 S_i d\lambda}\right) + 5\log_{10}\left(\frac{d}{r}\right). \quad (5)$$

Here F_0 is the zero-point flux density appropriate for each filter (the flux density that corresponds to a magnitude of zero). The new parameter d is introduced as the distance to the cluster. If we substitute Equations 4 into 5, we get

$$m'_i(t, q_{\rm H}, A_\nu, e, d). \quad (6)$$

At this point m' is used to describe the magnitude rather than m, as the photometric biases described in Section 3.1 have not been taken into account yet. The two filters used in the WFC3 observations are F225W and F336W. Magnitudes in these filters will be referred to as m_2 and m_3, respectively.

The thing we wish to calculate is the probability density, f, of finding a white dwarf at a given set of magnitudes (M_2 and M_3) and a given position in the field (R). This is $f = \frac{dN}{dm_2 dm_3 dR}$. In order to manipulate Equation 6 into the form we need to compare to the data, we first solve for t as a function of everything else, and then take two derivatives:

$$\frac{dt}{dm'_2 dm'_3} = \frac{dt}{d(m'_2)^2}\frac{dm'_2}{dm'_3}. \quad (7)$$

Next we multiply by the formation rate of objects in the field, $\dot{N} = \frac{dN}{dt}$:

$$\frac{dN}{dm'_2 dm'_3} = \frac{dt}{d(m'_2)^2}\frac{dm'_2}{dm'_3}\frac{dN}{dt}. \quad (8)$$

Now we have to bring back the error distribution, E, from Section 3.1. Before we refered to this as $E(\mathrm{F225W_{input}}, \mathrm{F336W_{input}}, R, \Delta\mathrm{F225W}, \Delta\mathrm{F336W})$, but now we can see that $\mathrm{F225W_{input}} = m'_2$, $\mathrm{F336W_{input}} = m'_3$, ΔF225W$= m_2 - m'_2$, and ΔF336W$= m_3 - m'_3$. So what we have is:

$$E(m_2 - m'_2, m_3 - m'_3, m'_2, m'_3, R). \quad (9)$$

We can perform an integral to transform Equation 8 into the form that we want, which is f:

$$f = \frac{dN}{dm_2 dm_3 dR} = \rho(R)\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\frac{dN}{dm'_2 dm'_3}E(m_2 - m'_2, m_3 - m'_3, m'_2, m'_3, R)dm'_2 dm'_3. \quad (10)$$

This step is a convolution integral of the error distribution with the theoretical two dimensional luminosity function. The error distribution changes with position in data-space. The assumed radial density distribution, $\rho(R)$, must also be included at this point. Neglecting this term is equivalent to an assumption of a density distribution that is constant with radius. We use a King-Michie model (King 1966; Michie 1963) to calculate the projected density distribution. The assumed King radius and tidal radius values are taken from Goldsbury et al. (2013) and Harris (1996), respectively ($r_0 = 32''$ and $r_t = 3800''$). Note that $\rho(R)$ in Equation 10 is not the three dimensional density distribution, it is the projected and azimuthally integrated distribution with units of objects per pixel.

At this point we have a distribution f that gives the probability density of finding an object at any point in the data-space ($F225W, F336W, R$) for any combination of our model parameters ($q_{\rm H}, A_\nu, d, e, \dot{N}$). We will use a test statistic referred to as the "unbinned likelihood" to determine how well the data are predicted by this model.

3.4. *The Unbinned Likelihood*

A quick derivation of the unbinned likelihood starts by considering a bin in the data-space with dimensions ΔR by Δm_2 by Δm_3. This bin is centred at at the location (R_j, m_{2k}, m_{3l}). From the model discussed above, the predicted number of objects in this bin is given as $f(R_j, m_{2k}, m_{3l})\Delta R \Delta m_2 \Delta m_3$. The probability of finding n objects in this bin is given by the Poisson distribution:

$$P(n; f(R_j, m_{2k}, m_{3l})) = \frac{[f(R_j, m_{2k}, m_{3l})\Delta R\Delta m_2\Delta m_3]^n e^{-f(R_j,m_{2k},m_{3l})\Delta R\Delta m_2\Delta m_3}}{n!}. \quad (11)$$

The limit is then taken as the bin size approaches zero. This leaves a number of bins equal to the number of data points with one object in them, and an infinite number of bins with zero objects in them. Then there are two possible forms for the probability to take:

$$P(1; f(R_j, m_{2k}, m_{3l})) = f(R_j, m_{2k}, m_{3l})\Delta R\Delta m_2\Delta m_3 e^{-f(R_j,m_{2k},m_{3l})\Delta R\Delta m_2\Delta m_3}, \quad (12a)$$

$$P(0; f(R_j, m_{2k}, m_{3l})) = e^{-f(R_j,m_{2k},m_{3l})\Delta R\Delta m_2\Delta m_3}. \quad (12b)$$

The log-likelihood is the sum of the logarithm of the probabilities of all of the bins:

$$\log L = \sum_i \log(f(R_i, m_{2i}, m_{3i})) + \sum_i \log(\Delta R\Delta m_2\Delta m_3) - \sum_{jkl} f(R_j, m_{2k}, m_{3l})\Delta R\Delta m_2\Delta m_3. \quad (13)$$

The first two sums are only performed over i, which is the index of the data points, since these terms only exist in the case of Equation 12a. The third sum is carried out for all bins, even those with no data in them, since this term is common to Equations 12a and 12b. The second sum does not depend in any way on the data or our chosen model. It will be an additive constant to the log-likelihood and so can be ignored. Given the limit as the bins become infinitesimal, we can also see that the third sum becomes an integral of the probability density over the entire data-space. With these two modifications we have

$$\log L = \sum_i \log(f(R_i, m_{2i}, m_{3i})) - \iiint_{\mathrm{data-space}} f dR dm_2 dm_3, \quad (14)$$

$$\log L = \sum_i \log(f(R_i, m_{2i}, m_{3i})) - N_{\mathrm{pred}}. \quad (15)$$

The integral in Equation 14 is the total number of objects that our model predicts in the region of data-space in which we are doing the comparison. To simplify the notation, we have written this as N_{pred} in Equation 15. The bound region in data-space does not have to be rectangular or any other particular shape; as long as the integral in Equation 14 is performed over exactly the same region from which we select the data, this works.

The description up until this point has been specific to the central WFC3 data, but a virtually identical approach is used for the ACS data. All nine ACS fields were reduced separately. Since we assume that the cluster distribution and also the incompleteness are azimuthally symmetric, we only perform artificial star tests on one of these nine images. These tests are still parametrized by radial position R and two magnitudes (in this case $F435W$ and $F555W$). We can use these tests to build a second probability density function f specific to the ACS data.

Although it is left out of the likelihood equations above, the probability density functions f, and therefore the likelihood as well, both depend on the choice of model parameters $(q_{\mathrm{H}}, A_\nu, d, e, \dot{N})$. The first four parameters are consistent across both the ACS data and the WFC3 data. This is because all of the stars should have the same hydrogen layer thickness, have the same neutrino physics, lie at the same distance, and experience the same reddening. The fifth parameter $\dot{N}$ is the rate at which white dwarfs are being produced in each field, so will not be the same for both data sets.

To understand the $\dot{N}$ parameter more intuitively, imagine first selecting the region shown in Figure 6. If we could watch this CMD change over hundreds of millions of years, we would see objects moving down the white dwarf sequence as they cool and leaving the region through the bottom. We would also see new white dwarfs arriving through the top of the region as they contract and cool coming from the horizontal branch. $\dot{N}$ would be the flux into the fit region as observed in this way. This flux changes on long timescales, as the cluster turn-off moves down the mass function to lower mass stars, but it will change slowly enough over the age range of white dwarfs in the fit region that we can assume that it is constant.

In total we now have six parameters; $(q_{\mathrm{H}}, A_\nu, d, e, \dot{N}_{\mathrm{WFC3}}, \dot{N}_{\mathrm{ACS}})$. We can combine the two likelihood functions to generate a likelihood over parameter space for all of the data together.

3.5. *Degenerate Parameters and Prior Knowledge*

A number of the parameter constraints are partially degenerate. Distance and reddening are partially degenerate, due to the fact that a larger reddening requires a larger extinction in each filter, producing a similar effect to putting the cluster farther away. The formation rate is degenerate with distance, which is because the cooling rate roughly follows a power law with time. For a perfect power law it is impossible to distinguish between a vertical and horizontal shift in log-space. Because the model cooling rate is similar, changing the timescale looks roughly the same as making all of the objects brighter or fainter.

In addition to these degeneracies, the data set is not particularly useful for obtaining constraints on these parameters. Without using any prior distributions on other parameters, the fit distributions for distance or reddening are 2-3 times wider than current constraints. This isn't worrying however, since the parameters we are most interested in studying are q_{H} and A_ν, and the data are able to provide useful constraints on these parameters. So for these reasons, we choose to use priors for all of our model

parameters except for $q_{\rm H}$ and A_ν. This leads to a constraint on the neutrino production rate and the mass of the hydrogen layer, dependent on what we already know about the other cluster parameters.

Table 2
Values for $\dot{N}$ are determined as described in the paragraph following this table. They are provided in units of objects per Myr.

Parameter	Mean	Std. Dev.	Reference
$(m-M)_0$	13.30	0.13	Woodley et al. (2012)
$E(B-V)$	0.04	0.02	Harris (1996)
$\dot{N}_{\rm WFC3}$ ($\rm Myr^{-1}$)	8.2	0.3	
$\dot{N}_{\rm ACS}$ ($\rm Myr^{-1}$)	2.61	0.07	

We use Gaussians for all of the prior distributions. The parameters of these are shown in Table 2. The values for the formation rate of white dwarfs in each field ($\dot{N}$) are estimated in the same way as in Goldsbury et al. (2012). We use data on the giant branch in each field, along with completeness corrections from the artificial star tests, and main sequence evolution models from Dotter et al. (2008) to model the rate at which stars are leaving the main sequence and evolving up the giant branch. We make the assumption that this is the same rate at which stars are arriving onto the white dwarf sequence.

3.6. *Restricting the Fitting Region*

The region in magnitude space over which we fit our model is restricted in both fields to the region immediately around the white dwarf sequence. These regions are shown in blue in Figure 6. There are two main reasons for the requirement that the fitting be restricted to these regions. In both the WFC3 and ACS data, the lower white dwarf sequence begins to overlap with the main sequence of the Small Magellanic Cloud, which lies in the background of 47 Tuc.

We do not include the SMC in our model and so fitting over this region would be unreliable. In the case of the WFC3 data, the white dwarf models also deviate significantly from the measured sequence at fainter magnitudes. The cause of this discrepancy is unclear. After a thorough analysis, we find no apparent systematic offset introduced by the photometric reduction process. The photometry of the sequence is slightly biased by diffuse background light and crowding of the inner field, but this could account for only 10% of the discrepancy observed at most. Additionally, the error model described in Section 3.1 fully accounts for the discrepancy attributable to diffuse background light. The model offset even persists through entirely different reductions performed with and without drizzling. Avoiding the lower WFC3 sequence does have a considerable impact on the total sample size, but if the fitting region were not restricted as shown in Figure 6, all of the white dwarf models would be strongly ruled out. Using the region shown, there are regions of parameter space that are consistent with the data. Although the plot shows this region in colour-magnitude space, fitting is done in magnitude-magnitude space.

3.7. *Markov Chain Monte Carlo Sampling*

To explore the model parameter space, we use a Markov chain Monte Carlo (MCMC) method. Specifically, we use the Metropolis-Hastings algorithm (Metropolis et al. 1953; Hastings 1970) with Gibbs sampling (Geman & Geman 1984). The total chain used to generate the parameter distributions in this analysis is 100,000 points long. The process of thinning, in which only every Nth point in a Markov chain is kept, is often used to reduce the correlation between points within the chain. We do not follow this approach, since these small-scale correlations are averaged out by the comparably long length of the total chain. This is explained in detail in Link & Eaton (2012). We can look at the autocorrelation of each of our model parameters to see that beyond 20 steps the auto-correlation drops to effectively zero (Figure 7). It is possible to break the total chain into many smaller chains, each of which gives the same result within statistical uncertainties (Figure 8). These simple tests indicate that the total chain will be not be biased by the small length scale correlations within it.

4. RESULTS

4.1. *Hydrogen Layer Thickness ($q_{\rm H}$) and Neutrino scaling (A_ν)*

The only two parameters fit without priors are A_ν and $q_{\rm H}$. The join fit distribution for these parameters is shown in Figure 9. The distribution for $q_{\rm H}$ alone is shown in Figure 10. This is the histogram of the $q_{\rm H}$ values in the total chain. It is also the likelihood distribution of this parameter after marginalizing over all of the other parameters. The range of 2.2×10^{-5} to 8.9×10^{-5} contains 95% of the distribution. The distribution for A_ν alone is shown in Figure 11. This is the histogram of A_ν values in the total chain. It is also the likelihood distribution of this parameter after marginalizing over all other parameters. The range of 0.83 to 1.22 contains 95% of the distribution. The default amount of neutrino production in the MESA models corresponds to $A_\nu = 1$. Our findings are consistent with this.

Figures 12 and 13 show how the cooling models change with varying A_ν and $q_{\rm H}$ respectively. From these figures it is easy to see that changing A_ν has a much larger effect than changing $q_{\rm H}$ over the range of values we are considering. This results in a tighter constraint on A_ν.

4.2. *Neutrino Species*

The MESA software uses the results of Itoh et al. (1996) to calculate the neutrino emissivity. The neutrino production in white dwarfs in this temperature range is dominated by the plasmon neutrino process. The term "plasmon" refers to a photon moving through a plasma. Unlike photons in a vacuum, plasmons travel more slowly than c due to their interaction with the plasma. Some of the energy of the photon is transferred to the free electrons in the plasma. This allows the interacting plasmon to decay into a neutrino-antineutrino pair while still conserving momentum: $\gamma \rightarrow \nu + \bar{\nu}$ (Kantor & Gusakov 2007). The factor we refer to as A_ν is a multiplication of $Q_{\rm plasma}$ from equation 4.1 of Itoh et al. (1996). The plasmon neutrino rate is calculated in the MESA code as described in Itoh et al. (1996), and this rate depends on the central density of the white dwarf, core temperature of the white dwarf, the weak mixing angle (more often parametrized as $\sin^2\theta_{\rm W}$), and the effective number of non-electron neutrino species (n). We multiply

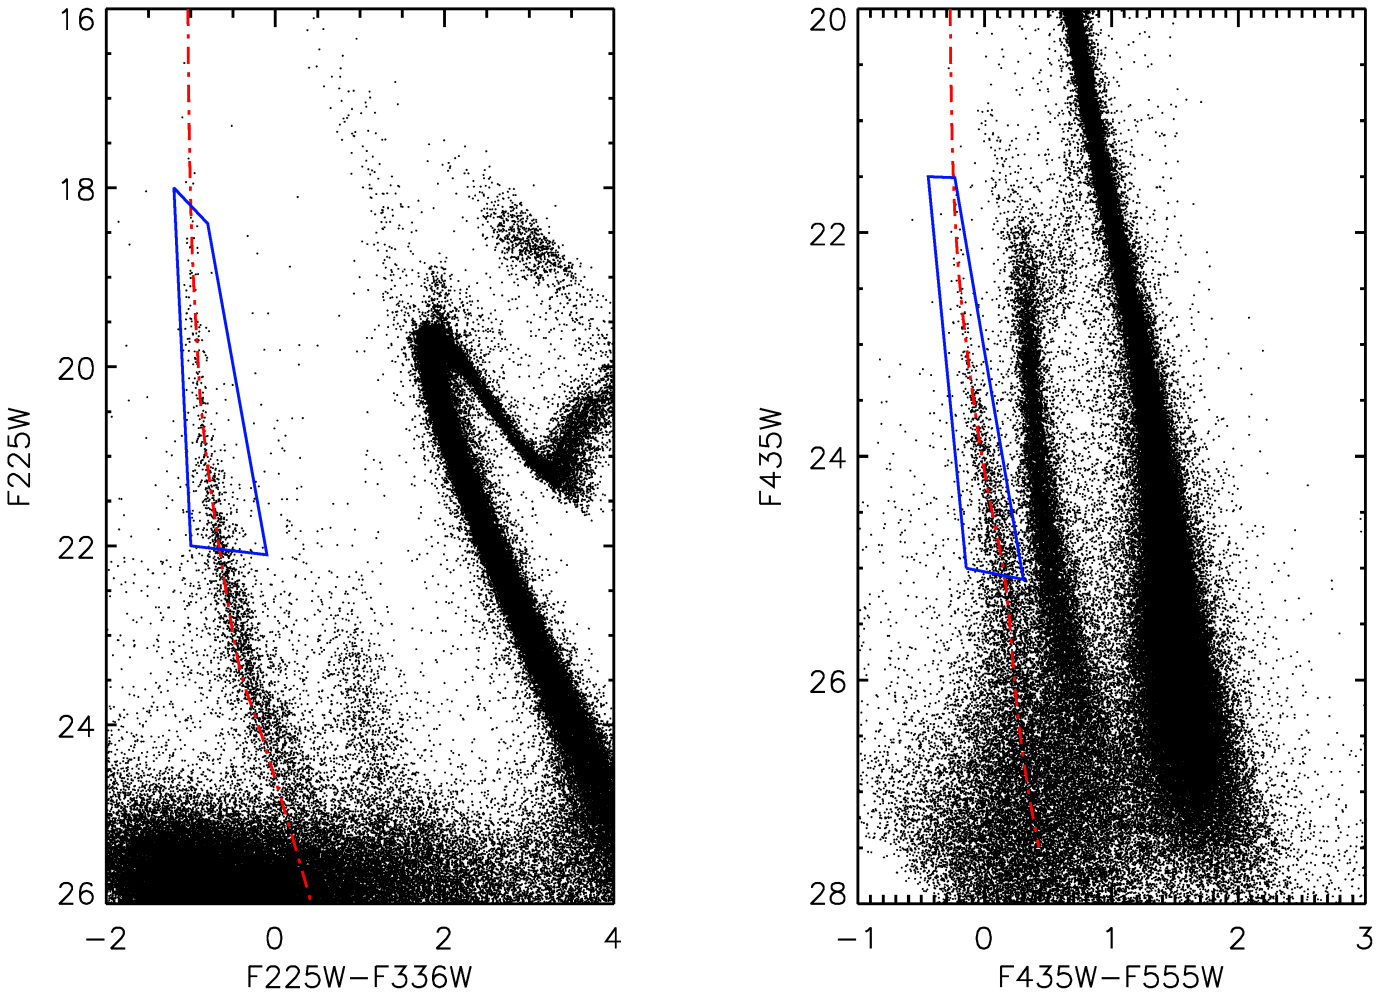

Figure 6. The colour magnitude diagrams of the two fields. Overplotted are the region in which we fit our model in blue, and the white dwarf cooling sequence in red for the mean prior values given in Table 2, before convolution with the error model.

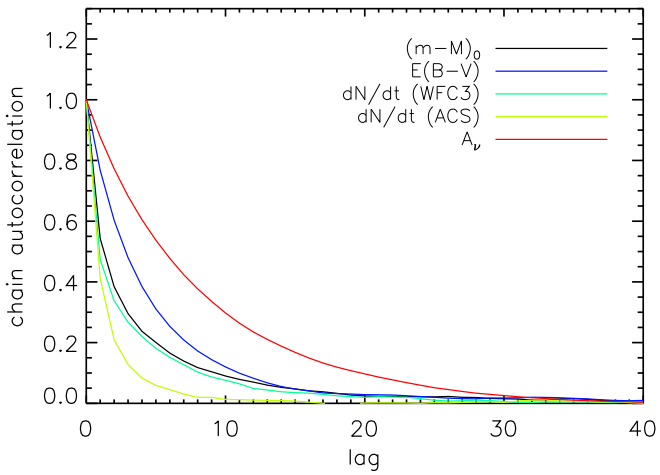

Figure 7. The auto-correlation of each chain parameter as a function of the lag between points in the chain. No memory of the past position in parameter space is retained beyond 40 steps.

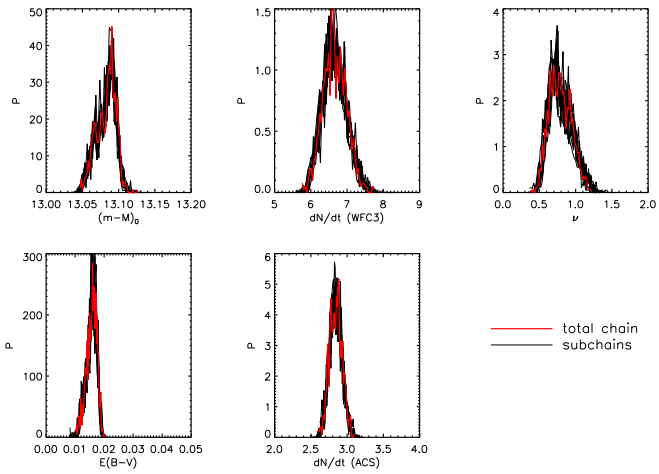

Figure 8. A comparison of the likelihood distribution for each of our model parameters shown for the entire chain in red. The thinner black lines are independent segments 1000 steps long. Each of these separate chains is much longer than the length scale on which correlations persist in the chains. All of them average to the same broad distribution that we find with the entire chain, but with more statistical noise.

the rate that comes out of this calculation by the factor A_ν. The rate $Q_{\rm plasma}$ can be affected by any of the four parameters listed above, and so the factor A_ν could be thought of as a change to any combination of these four. However, interpreting A_ν as a change to either the central density or temperature would be inconsistent, since these parameters were not consistently altered elsewhere in the code. It therefore only makes sense to interpret a constraint on this parameter as either a constraint on

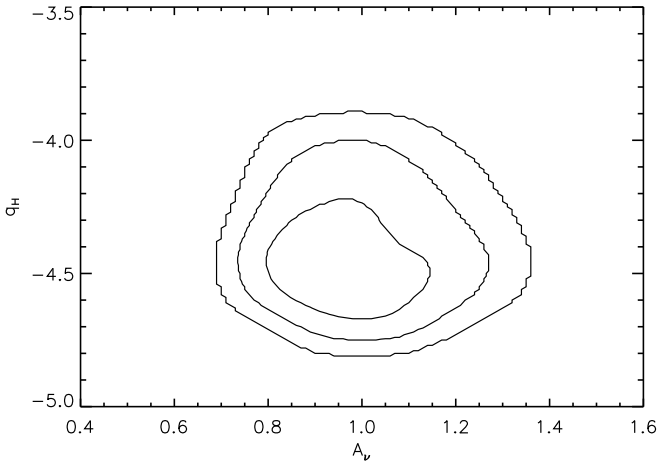

Figure 9. Contours of the join fit distribution of A_ν and $q_{\rm H}$. 1, 2, and 3σ contours (enclosing 68%, 95%, and 99.7%) are shown.

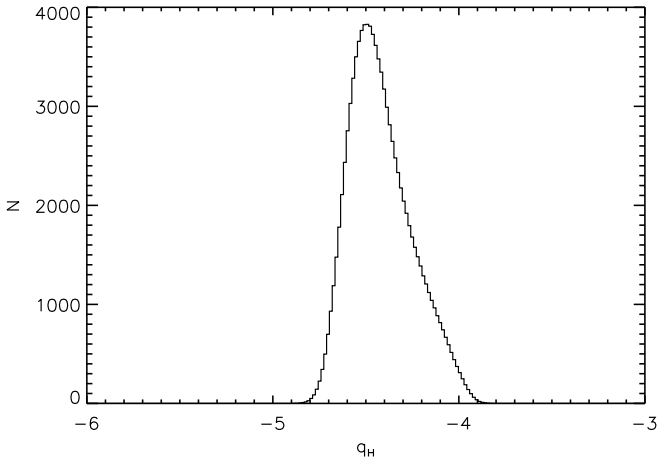

Figure 10. A histogram of the posterior distribution of $q_{\rm H}$ values from the markov chain. This is the likelihood distribution of $q_{\rm H}$ after marginalizing over all other parameters in the model with prior distributions on each, except for A_ν. The range of 2.2×10^{-5} to 8.9×10^{-5} contains 95% of the distribution.

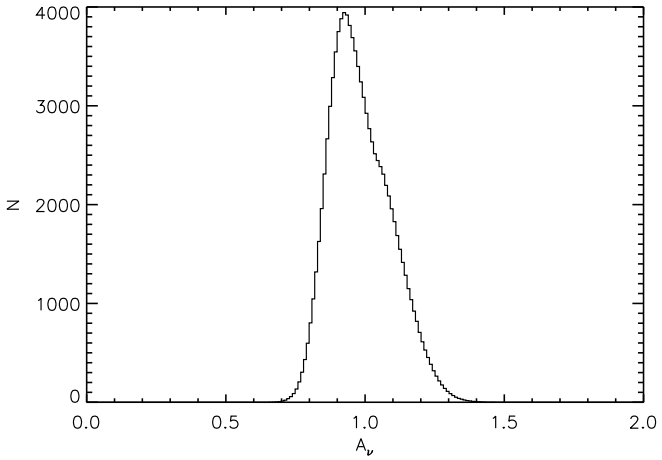

Figure 11. A histogram of the posterior distribution of A_ν factor values from the markov chain. This is the likelihood distribution of A_ν after marginalizing over all other parameters in the model with prior distributions on each, except $q_{\rm H}$. The range of 0.83 to 1.22 contains 95% of the A_ν fit distribution.

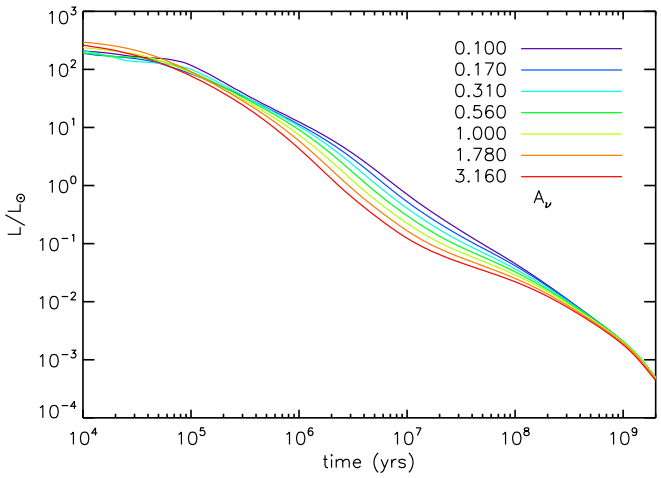

Figure 12. The relation between luminosity and time shown for varying A_ν values while $q_{\rm H}$ is held fixed.

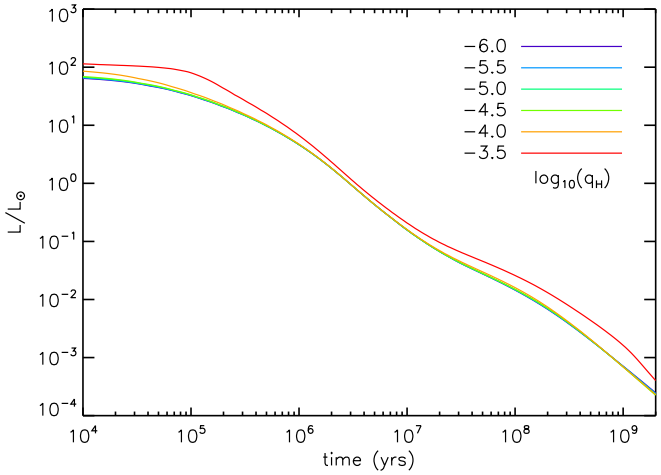

Figure 13. The relation between luminosity and time shown for varying $q_{\rm H}$ values while A_ν is held fixed.

$\sin^2\theta_{\rm W}$ and n.

The relation between A_ν, $\sin^2\theta_{\rm W}$, and n is

$$A_\nu \propto (C_V^2 + nC_V'^2), \tag{16}$$

where C_V and C_V' are defined as

$$C_V = \frac{1}{2} + 2\sin^2\theta_{\rm W} \tag{17}$$

and

$$C_V' = 1 - C_V. \tag{18}$$

Using these relations, the likelihood in Figure 11 can be mapped onto the two-dimensional space of $\sin^2\theta_{\rm W}$ and n. This is shown in Figure 14. From the Particle Data Group (Olive et al. 2014) $\sin^2\theta_{\rm W} = 0.23155$. Looking at this value in Figure 14, one can see that changing n will not produced a large change in A_ν. This means that our constraint on A_ν does not allow us to put a reasonable constraint on n.

4.3. *Fit Quality*

While these distributions in parameter space inform us which values lead to better or worse fits, they can-

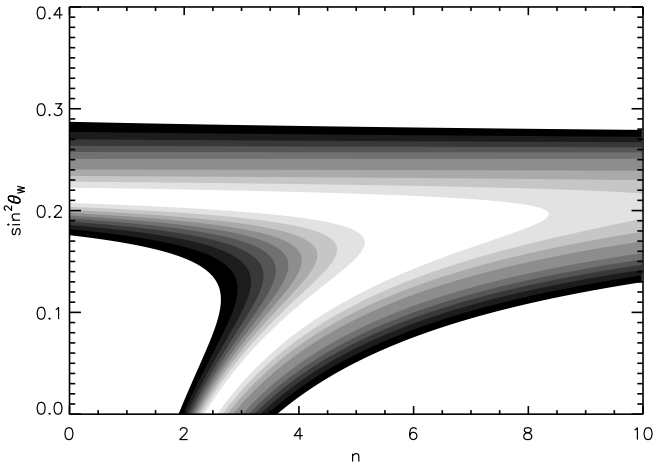

Figure 14. The likelihood of $\sin^2\theta_W$ and n. For the accepted value of $\sin^2\theta_W$ (0.23155 from Olive et al. (2014)), C'_V is very small. This means that the dependence of A_ν on n is very weak. As a result, our constraint on the neutrino production (A_ν) does not reasonably constrain n.

not indicate which fits are actually consistent with the dataset. This is possible in the same way that one could determine a best-fitting line for data that are clearly not linear. A maximum in likelihood will be found regardless of whether or not the fit is good. It is possible that even though we have found constrained likelihood distributions for q_H and A_ν, the best fitting model would still be a poor fit to the data.

To evaluate whether this is the case we use our best-fitting model to create simulated data through Monte Carlo sampling. The model being referred to here is Equation 10, so these simulated samples take full account of the empirically determined incompleteness as well as photometric bias and scatter. We generate millions of simulated data sets from the best-fitting model and in each case calculate the unbinned likelihood of the simulated data against the model. These values make up the distribution of likelihood we would expect if data such as ours were drawn from our best-fitting model. We can then ask whether or not the likelihood value we find from comparing the real data to the model is an outlier in this distribution. If the real likelihood value falls well outside the distribution of likelihood values from the simulated data, then we would conclude that it is unlikely the real data could be drawn from our best-fitting model. The distribution of likelihood values from simulated samples, as well as the likelihood value of the real data compared to the model, are shown in Figure 15.

Integrating from the red dashed line to the left tail of the distribution in Figure 15 will enclose 4.9% of the total area under the black line. This indicates that 4.9% of the time we would expect to draw data from our model that differs at least as much as the real data do.

4.4. *Cooling Models From Other Groups*

We compared the four cooling models used in Goldsbury et al. (2012) to this data set as well. These models come from Fontaine et al. (2001), Hansen et al. (2015), Lawlor & MacDonald (2006), and Renedo et al. (2010). The model from Hansen et al. (2015) is more recent than the version fit in Goldsbury et al. (2012), while the other models are the same. These models do not have any intrinsic parameters to be varied, so the fits are only

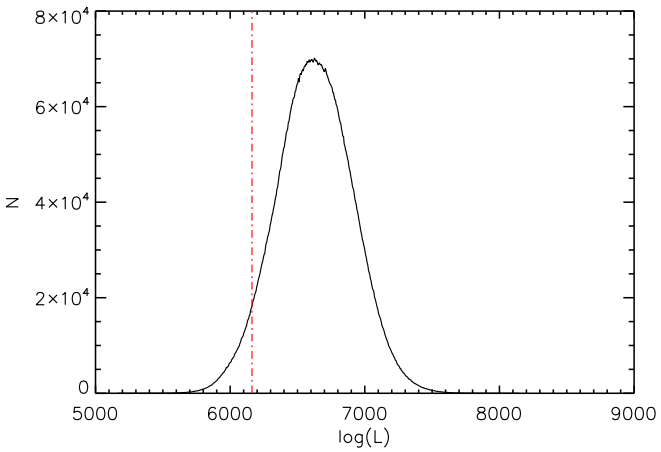

Figure 15. The distribution of likelihood values expected when sample data are drawn from the best fitting cooling model is shown as the black distribution. The red line indicates the log-likelihood value found when comparing the real data to the model. We would expect to draw data from our best fitting model that differs at least as much as the real data do from that model 4.9% of the time.

performed over the parameter space (with priors) defined by the four parameters in Table 2. The quality of fit assessment for each model is done in the same way as for our MESA models (summarized in Figure 15. Simulated data is repeatedly drawn from each best-fitting model and compared back to that model. The distribution of resulting likelihood values is then integrated from the value obtained when comparing the real data to the model. The results are shown in Table 3. With this data set we cannot strongly rule out any of these models.

Table 3
The "fit probability" is defined as the probability that data drawn from the given model would produce a likelihood value less than the real data do when compared back to that model.

cooling model	fit probability
Fontaine et al. (2001)	0.032
Hansen et al. (2015)	0.025
Lawlor & MacDonald (2006)	0.078
Renedo et al. (2010)	0.118
this work (MESA)	0.049

5. CONCLUSIONS

We have described the implementation of a fitting statistic, known as the unbinned maximum likelihood, which is common in some fields but relatively unused in astronomy. This approach does not require that the data be binned in any way, or an assumption of Gaussian error bars, both of which are necessary for a χ^2 fit. Given the same data and model, an unbinned approach will allow for tighter constraints on the same parameters than a binned fit will. In the case we have described, the error model is combined directly with the theoretical model to produce a continuous distribution for where points are expected to be found in data-space. The data can be kept unmodified from their raw state, and all assumptions about parameters and errors are built into the model itself.

We used this approach to test white dwarf cooling models against the largest sample of white dwarfs ever collected in a single globular cluster (47 Tuc). Fitting in

six-dimensional parameter space was performed with an MCMC sampler. We found constraints on the thickness of the Hydrogen layer and the amount of neutrino production in the white dwarf cores.

The constraint on $q_{\rm H}$ is shown in Figure 10. The data prefer thicker hydrogen layers ($q_{\rm H} = 3.2 \times 10^{-5}$) and we can strongly rule out thin layers ($q_{\rm H} = 10^{-6}$). The 95% range of this parameter is 2.2×10^{-5} to 8.9×10^{-5}. Our constraint on neutrino production is much better than for $q_{\rm H}$, with $0.83 < A_\nu < 1.22$ at 95% confidence. In Hansen et al. (2015) the best-fitting hydrogen layer thickness was found to be $q_{\rm H} = 4 \times 10^{-4}$. Our analysis strongly rules out that value, but we also find better fitting hydrogen layers at the thick end of the range usually considered by modellers (Bergeron et al. 2001). The 95% confidence range on f_s in Hansen et al. (2015) (which is equivalent to our parameter A_ν) was found to be $0.6 < f_s < 1.7$, which agrees with what we have found. The differences in these values can be attributed both to the inclusion of additional data in our analysis, and to different statistical assumptions made in fitting the models.

REFERENCES

Althaus, L. G., Camisassa, M. E., Bertolami, M. M. M., Córsico, A. H., & García-Berro, E. 2015, Astronomy & Astrophysics, 576, A9

Bergeron, P., Leggett, S., & Ruiz, M. T. 2001, The Astrophysical Journal Supplement Series, 133, 413

Blöcker, T. 1995, Astronomy and Astrophysics, 297, 727

Dotter, A., Chaboyer, B., Jevremović, D., et al. 2008, The Astrophysical Journal Supplement Series, 178, 89

Fitzpatrick, E. L. 1999, 26, 63

Fontaine, G., Brassard, P., & Bergeron, P. 2001, Publications of the Astronomical Society of the Pacific, 113, 409

Fruchter, A., Sosey, M., Hack, W., et al. 2009, STScI, Baltimore

Geman, S., & Geman, D. 1984, Pattern Analysis and Machine Intelligence, IEEE Transactions on, 721

Goldsbury, R., Heyl, J., & Richer, H. 2013, The Astrophysical Journal, 778, 57

Goldsbury, R., Heyl, J., Richer, H., et al. 2012, The Astrophysical Journal, 760, 78

Hansen, B., Richer, H., Kalirai, J., et al. 2015, arXiv preprint arXiv:1507.05665

Hansen, B., Kalirai, J., Anderson, J., et al. 2013, Nature, 500, 51

Hansen, B. M., Anderson, J., Brewer, J., et al. 2007, The Astrophysical Journal, 671, 380

Harris, W. E. 1996, AJ, 112, 1487

Hastings, W. K. 1970, Biometrika, 57, 97

Heyl, J., Richer, H. B., Antolini, E., et al. 2015, The Astrophysical Journal, 804, 53

Itoh, N., Hayashi, H., Nishikawa, A., & Kohyama, Y. 1996, The Astrophysical Journal Supplement Series, 102, 411

Kantor, E., & Gusakov, M. 2007, Monthly Notices of the Royal Astronomical Society, 381, 1702

King, I. R. 1966, AJ, 71, 64

Lawlor, T., & MacDonald, J. 2006, Monthly Notices of the Royal Astronomical Society, 371, 263

Link, W. A., & Eaton, M. J. 2012, Methods in Ecology and Evolution, 3, 112

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, The journal of chemical physics, 21, 1087

Michie, R. W. 1963, MNRAS, 126, 499

Olive, K. A., et al. 2014, Chin. Phys., C38, 090001

Paxton, B., Bildsten, L., Dotter, A., et al. 2011, The Astrophysical Journal Supplement Series, 192, 3

Reimers, D. 1975, Memoires of the Societe Royale des Sciences de Liege, 8, 369

Renedo, I., Althaus, L. G., Bertolami, M. M., et al. 2010, The Astrophysical Journal, 717, 183

Renzini, A., Bragaglia, A., Ferraro, F. R., et al. 1996, The Astrophysical Journal Letters, 465, L23

STScI. 2009, Multi-Cycle Treasury Programs

Tremblay, P. E., Bergeron, P., & Gianninas, A. 2011, 730, 128

Winget, D. E., Kepler, S. O., Campos, F., et al. 2009, The Astrophysical Journal Letters, 693, L6

Wood, M. 1992, The Astrophysical Journal, 386, 539

Woodley, K. A., Goldsbury, R., Kalirai, J., et al. 2012, The Astronomical Journal, 143, 50

Zoccali, M., Renzini, A., Ortolani, S., et al. 2001, The Astrophysical Journal, 553, 733